

TECHNICAL RESEARCH REPORT

Negotiating Access Control Policies Between Autonomous Domains

by Vijay G. Bharadwaj, John S. Baras

CSHCN TR 2002-3
(ISR TR 2002-3)



The Center for Satellite and Hybrid Communication Networks is a NASA-sponsored Commercial Space Center also supported by the Department of Defense (DOD), industry, the State of Maryland, the University of Maryland and the Institute for Systems Research. This document is a technical report in the CSHCN series originating at the University of Maryland.

Web site <http://www.isr.umd.edu/CSHCN/>

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2002		2. REPORT TYPE		3. DATES COVERED -	
4. TITLE AND SUBTITLE Negotiating Access Control policies Between Autonomous Domains				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency,3701 North Fairfax Drive,Arlington,VA,22203-1714				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Negotiating Access Control Policies Between Autonomous Domains

Vijay G. Bharadwaj and John S. Baras

*Institute for Systems Research, University of Maryland,
College Park MD 20742, USA.*

{vgb,baras}@umd.edu

Abstract

Autonomous policy domains often need to share resources to accomplish a common task. To do this they must negotiate a common access control policy to the shared resources. We use mathematical techniques from game theory to show that the outcome of such negotiations can often be predicted from the distribution of power among the participants, independent of the actual mechanics of negotiation. We discuss the axiomatic derivation of some game theoretic solution concepts, and illustrate our techniques with examples.

1. Introduction

Online collaboration frequently requires two or more autonomous policy domains to form a coalition in order to share data or other resources to achieve a common goal. Often, the collaboration itself may generate new data or resources, and these must also be shared. Therefore the domains involved must negotiate a common access control policy for all shared resources. Traditionally, such negotiations have been carried out by human beings meeting in person, through a tedious process of discussion and bargaining that can take weeks or months to conclude. In many situations, especially when the collaboration is for a short period of time, these delays are unacceptable. It is then desirable to speed up the process of negotiation by automating it.

Previous work on negotiation either assumes that coalition policies are agreed upon by extra-technological means [1], or treats the problem in the context of client-server networks [2]. We look at networks with multiple peer domains, all of whom share resources with each other for mutual benefit. Our aim is to automate the process of policy negotiation in such networks, so that it can be carried out by software agents (perhaps residing on the domain controllers) with limited human intervention.

Conceptually, when negotiating a common access control policy, the domains must agree on two things: a model and a set of rules for the mechanisms by which resources

are shared (i.e. the policy model and the policy model interpretation), and which resources are shared and with whom (i.e. the access state of the coalition under the access control policy). These two parts are related: for instance, the access state must be compatible with the properties of the agreed-upon policy model.

Previous work [3] has explored policy definition languages for security domains; such languages can also be extended to describe coalition policies. For instance, consider a coalition in which all domains use Role Based Access Control as their policy model. Each domain can share its resources with foreign domains by enrolling some users from the foreign domain into local roles with access to the relevant resources. Another approach is to create a new set of roles for the coalition, then give these new roles access to the shared resources and enroll local and foreign users into these new roles. In either case, languages for describing role-based access control models can be used to express coalition policies as well. Thus, domains in a negotiation can communicate their security policies to each other, and can automatically check if a proposed state is consistent with the policy.

In this paper we look at the negotiation process itself. We start by postulating some desirable properties we would like the negotiated coalition policy to have. We then show that these properties lead naturally to certain game theoretic concepts, and to algorithms for computing them. As a result we can characterize the likely outcomes of policy negotiation and provide methods to compute these outcomes.

Our results show that the outcome of a negotiation is often determined by the distribution of power among the participants rather than by the details of how the negotiation proceeds. The power of a domain in a negotiation is its ability to help or hurt the other domains by cooperating with them or refusing to do so. The language of game theory allows us to frame these properties in a precise manner, and gives us the mathematical tools to predict the eventual results of such negotiations.

We assume an architecture in which each domain has a central controller to administer its local policies and to negotiate with other domains on its behalf (see Figure 1).




Figure 1: Negotiation in multi-domain networks.

We will refer to such networks as coalitions, and to the resources shared between the domains in the network as coalition resources.

The rest of this paper is organized as follows: Section 2 introduces some basics of game theory, Section 3 introduces the examples we will use to illustrate our techniques, and Sections 4 and 5 show some tools to analyze the examples and predict their outcomes. Section 6 concludes the paper.

2. Game Theory

Game theory [4, 5] is the mathematical study of conflict and cooperation between intelligent rational entities (referred to as players). By modeling such situations mathematically, game theory can predict what kinds of cooperation will arise in a group of players under a given set of conditions. The same tools can also be applied to design games that lead to desirable outcomes; we can decide what kinds of cooperation we would like to see in a group of players, and devise rules for their interaction such that it will be in the players' best interests to cooperate with each other.

A game theoretic model consists of a set of players, a set of possible actions for each player, and a payoff function, which associates each combination of actions by the players to a vector of rewards to the players. A key assumption is that all players are rational and intelligent. Rationality means that each player behaves in his own best interest, i.e. in a manner calculated to increase his own reward. Intelligence implies that players are capable of making any inferences about the structure and dynamics of the game that we are. Thus players are capable of deducing what their optimal strategies are in any given situation, as well as what the

optimal strategies of the other players are.

A solution concept is a rule that associates a game with a set of outcomes. Many solution concepts have been proposed in the literature. Each solution concept is based on a set of axiomatic requirements that the solution must satisfy, and allows us to make predictions about some aspect of player behavior in a given game. Most useful solution concepts also provide algorithms to compute solutions that satisfy their axioms for any given game.

Noncooperative games are those in which every player tries to increase her own reward independently of all the other players. Cooperative games, on the other hand, are those in which players can organize themselves into groups to achieve a common goal that benefits them all. Often, cooperative games include the concept of transferable utility - there exists a commodity, such as money, which can be freely transferred among the players and which serves as a common standard of value.

Due to the large variety of possible cooperation structures, cooperative games are often richer in structure and offer deeper insight into many real-life situations. In this paper we use solution concepts from both noncooperative game theory and cooperative game theory to reason about the problem of automated policy negotiation.

As we will show in this paper, policy negotiations can be cast as games, where the outcomes of the game are the various policies that may be agreed on through negotiation. The dynamics of the game depend on the rules of negotiation, and we will disregard these, assuming only that they are flexible enough to allow the participants sufficient opportunity to communicate with each other and arrive at a result. We also assume that unanimous agreement of all participants is required to terminate a negotiation. We will seek solutions of these games, i.e. we will try to make predictions about how the negotiations will end, based on assumptions about what kinds of negotiated outcomes are desirable.

3. Two Examples

We now describe the two examples we will use in this paper to illustrate the use of game theory in analyzing policy negotiation. In both these examples we assume there are no external rewards or costs associated with the problem; that is, there is no compulsion on the negotiators to reach agreement and nothing to be gained from an agreement except for the benefits provided by the agreement itself. This assumption does not always hold true in applications; in many situations, the decision to collaborate has already been made, and negotiators are constrained by their need to reach agreement. We assume an absence of exogenous rewards or costs for simplicity; any quantifiable exogenous rewards or costs




Figure 2: Network for bandwidth sharing example.

can be easily incorporated into our models and dealt with using the same mathematical tools.

3.1. Bandwidth Sharing

Consider the network shown in Figure 2. There are two cities X and Y, and three Internet providers (labeled D1, D2 and D3) who wish to send data between them. All three Internet providers have many customers, with large amounts of data to send, and so would like to obtain as much capacity as possible. However, no provider owns a communication link that goes all the way from X to Y. Instead they own links that interconnect X and Y with other cities, labeled A through F. If the providers pool their resources, they can have a network capable of sending data from X to Y. They wish to negotiate a policy for sharing the transmission capacity between X and Y created by such a collaboration.

A number of features of this problem are apparent from the figure:

- None of the providers can send any data from X to Y on its own.
- No data can be sent from X to Y without D1's cooperation.
- D1 and D2 together can achieve a total rate of 1 Mbps.
- D1 and D3 together can achieve a total rate of 0.5 Mbps.
- All the providers together can achieve a combined rate of 3.5 Mbps.

We will denote by B_1 , B_2 and B_3 the share of the aggregate data capacity assigned to D1, D2 and D3 respectively. In later sections we discuss the values for B_1 , B_2 and B_3 that are likely to be produced by a negotiation between the providers.

3.2. Intelligence Sharing

In this example, three intelligence agencies (labeled A1 through A3) wish to share intelligence on targets of their interest. Each agency has a certain number of sources, and

Table 1: Parameters for intelligence sharing example.

Agency	No. of sources	Probability of compromise
A1	7	0.1
A2	4	0.5
A3	10	0.2

can give the other agencies access to as many of its sources as it wants. However, sharing a source increases the probability that it will be compromised. Each agency has a probability of compromise p_i , which represents the danger that a source accessible to this agency will be compromised. This danger applies equally to all sources a source knows about, including those it owns. Intelligence from a compromised source is useless, but intelligence from all uncompromised sources is equally valuable. The probability of a source being compromised due to one of the agencies is assumed to be independent of all other compromises. Table 1 shows the initial distribution of sources and the compromise probabilities. Assume that the compromise of a source cannot be detected by the agencies.

Denote by n_i the number of sources owned by the i th agency, and let s_i be the number of sources that the i th agency shares with the others (a source that is shared must be shared with all the other agencies). Let v_i , the reward to the i th agency (also known as the value derived by the i th agency, or the worth of the i th agency), be equal to the expected number of uncompromised sources that the agency has access to. Then, if no sharing is in effect,

$$v_i = n_i(1 - p_i)$$

If sources are shared, then v_i depends on the additional probability of compromise due to the other domains and the number of sources available from the other domains. Thus

$$\begin{aligned} v_i &= (n_i - s_i)(1 - p_i) + \left(\sum_j s_j\right) \prod_l (1 - p_l) \\ &= n_i(1 - p_i) + \left(\sum_j s_j\right) \prod_l (1 - p_l) - s_i(1 - p_i) \end{aligned}$$

Therefore each agency derives a benefit, in the form of an additional reward, from intelligence sharing. However, this benefit may in fact be negative if the agency exposes too many of its own objects to a threat of compromise. In what follows we show how game theoretic techniques can be used to predict which sharing arrangements are likely to arise.

4. Noncooperative Game Analysis

A fundamental concept in game theory is that of individual rationality. This is the simple idea that no player will voluntarily agree to an arrangement that makes him worse

off than he was before. It seems logical to suppose that no domain in a negotiation would agree to a policy that was not individually rational for that domain.

In practice, individual rationality allows us to reduce the set of possible outcomes to consider in a negotiation, by eliminating those outcomes that could never arise. A sharing arrangement is individually rational for a player if the value derived by that player under the arrangement is greater than the worth of the player without the arrangement.

In our bandwidth sharing example, none of the providers could send any data from X to Y on its own. The ability to send any data at all would be regarded as an improvement by each of the domains. Thus the individually rational solutions would only have to satisfy

$$B_1 \geq 0, \quad B_2 \geq 0, \quad B_3 \geq 0$$

However, in the intelligence sharing example, the constraints are not so trivial. Simple algebra shows that individually rational solutions must satisfy

$$\begin{aligned} \frac{s_1}{s_1 + s_2 + s_3} &\leq 0.4 \\ \frac{s_2}{s_1 + s_2 + s_3} &\leq 0.72 \\ \frac{s_3}{s_1 + s_2 + s_3} &\leq 0.45 \end{aligned}$$

As expected, Agency A2 needs to share more than the others, since its high probability of compromise makes its sources less trustworthy. On the other hand, Agency A1, which has the lowest probability of compromise, runs a significant risk by sharing its sources, and so has less incentive to share.

Individual rationality can help in the analysis of negotiation by weeding out irrational alternatives. This can identify games in which no individually rational solution exists, which would mean that no negotiation would succeed without the introduction of some outside reward or compensation. However, many games (like our examples above) have large numbers of individually rational solutions, and in these cases we need a rule to pick one out of these solutions.

Nash [6] showed that under mild technical conditions any game has a unique solution, known as the Nash bargaining solution, which satisfies the following axioms.

- *Individual rationality*: The solution is individually rational for all players.
- *Symmetry*: If two players have identical resources and identical reward functions, then they will receive equal treatment.
- *Scale Covariance*: Scaling the reward functions of the players by transformations of the form $\alpha x + \beta$ where

$\alpha > 0$ (α and β can be chosen differently for different players) does not affect the solution except to scale it by a similar transformation.

- *Pareto optimality*: No other solution exists which makes all the players better off than under this solution.
- *Independence from unfavorable alternatives*: Introducing a number of inferior outcomes into the structure of the game does not affect the solution.

Nash showed that the above axioms are satisfied by the individually rational solution which maximizes the product of the gains made by each player as a result of coalition formation. That is, if u_i was the worth of the i th player before coalition formation and x_i is his value after the coalition forms, the Nash solution is given by

$$\max_{\text{all possible outcomes}} \prod_i (x_i - v_i)$$

subject to

$$x_i \geq v_i \quad \forall i$$

For the bandwidth sharing example, the Nash bargaining solution is obtained by maximizing the product $B_1 B_2 B_3$ subject to $B_1 + B_2 + B_3 = 3.5 \text{ Mbps}$, which gives

$$B_1 = B_2 = B_3 = \frac{3.5}{3} \text{ Mbps}$$

The equal division reflects the fact that any provider which unilaterally breaks away from the coalition will not be able to send any data between X and Y, and so in a sense all the providers need the coalition equally.

For the intelligence sharing example, the Nash bargaining solution is obtained by maximizing

$$(s' - 0.9s_1)(s' - 0.5s_2)(s' - 0.8s_3)$$

where

$$s' = 0.36(s_1 + s_2 + s_3)$$

subject to

$$s_1 \leq 7, \quad s_2 \leq 4, \quad s_3 \leq 10$$

which gives us

$$s_1 = 4, \quad s_2 = 4, \quad s_3 = 5$$

Thus A2, which has the highest probability of compromise, must share as many sources as it can, and the other agencies will not share all their sources with A2.

5. Cooperative Games

The analysis in Section 4 has one shortcoming: it assumes that if negotiations between the players fail, then no coalition is formed at all. However, in real life, it is always possible for some of the players to negotiate with each other and form a smaller coalition, leaving out some players, if they find that this is more favorable to them. Thus, instead of choosing between forming the largest possible coalition and no coalition at all, players can choose between all the possible subsets of the set of players to form a coalition. Cooperative game theory is the study of games where the possibility for such coalition formation exists.

A cooperative game consists of a set of players, and a characteristic function that assigns a value to each subset of this set of players. The set of players is known as the grand coalition, and its subsets are known as coalitions. The value of a coalition is interpreted as the payoff, or total reward, that the players in that coalition can achieve by cooperating with each other. The game is known as a TU game (or game with Transferable Utility) if the players in a coalition are free to distribute their payoff among themselves in any way they choose; if players do not have such control over the distribution of payoffs, we have a game with Nontransferable Utility (or NTU game).

The best-known solution concept for cooperative games is the core. The core is the set of solutions that give to every possible coalition at least as much payoff as that coalition could get if it acted without the support of the other players. Thus the core is a measure of stability; if the players in a game agree to split their profits according to a core distribution, then no player or set of players will find it in their interest to secede from the grand coalition.

Before we define the cores for our examples, we must define the characteristic function for each of our examples. In other words, we must define what we mean by the worth of a coalition (as opposed to the worth of a single player). For the bandwidth sharing example, this is easy. We define the worth of any coalition as the maximum transmission capacity the members of that coalition can obtain without help from the remaining member(s).

For the intelligence sharing example, the value of a coalition is harder to define. Defining the value as the expected number of uncompromised sources available to the coalition is not quite satisfactory, as it disregards any sharing of sources that may take place. Therefore we define the value of any coalition as the sum of the values in the Nash bargaining solution for that coalition. For example, the value of the coalition {A1,A2} would be the sum of the values of A1 and A2 when they implement the Nash bargaining solution for the coalition consisting of A1 and A2 only. We contend that this is a reasonable definition, because if A1

and A2 are to share any of their sources with A3, they will only do so if they can get a better deal for themselves than they could without A3.

With these definitions, the core for the bandwidth sharing example is the (fairly large) set of bandwidth allocations that satisfy the inequalities

$$\begin{aligned} B_1 &\geq 0 \\ B_2 &\geq 0 \\ B_3 &\geq 0 \\ B_1 + B_2 &\geq 1 \\ B_1 + B_3 &\geq 0.5 \\ B_1 + B_2 + B_3 &= 3.5 \end{aligned}$$

For the intelligence sharing example, the core is characterized by

$$\begin{aligned} V_1 &\geq 6.3 \\ V_2 &\geq 2 \\ V_3 &\geq 8 \\ V_1 + V_2 &\geq 9.9 \\ V_1 + V_3 &\geq 24.48 \\ V_2 + V_3 &\geq 11.2 \end{aligned}$$

It turns out that there is no intelligence sharing arrangement (i.e. no values of $s_1 \leq 7$, $s_2 \leq 4$ and $s_3 \leq 10$) for which all the above inequalities hold. In other words the core of the intelligence sharing example is empty.

As seen from our examples, the core of a game may either be empty or so large as to be of little value in predicting player behavior. In the latter case, it is reasonable to expect that players will look for a core allocation that is the most efficient in some sense. Two solution concepts which emphasize fairness and efficiency are the nucleolus and the Shapley value. They are both unique and defined for large classes of cooperative games.

The nucleolus [7] is the solution that distributes payoffs in such a way as to lexicographically minimize the dissatisfactions of all the coalitions in the grand coalition. The dissatisfaction of a coalition is the reduction in its payoff as a result of joining the grand coalition. Equivalently, the nucleolus maximizes the benefit to the least rewarded member of the coalition. Thus the nucleolus tries to be fair to all the players, and is likely to be the outcome of a negotiation where the participants consider fairness to be important. The nucleolus can be shown to lie in the core if the core is nonempty. The nucleolus is also the unique solution having the following properties:

- *Anonymity*: The solution is independent of the labeling of the players.
- *Scale Covariance*: Scaling the reward functions of the players by transformations of the form $\alpha x + \beta$ where

$\alpha > 0$ (α and β can be chosen differently for different players) does not affect the solution except to scale it by a similar transformation.

- *Imputation Saving Reduced Game Property*: Loosely speaking, this means that if a subset of the grand coalition is prevented from participating in the negotiation but is still given a choice of whether to cooperate or not, the payoff to the remaining players is not affected.

The problem of finding the nucleolus for the bandwidth sharing problem can be reduced to a linear programming problem. Since it is a lexicographic minimization, in the worst case we have to solve as many linear programs as there are possible coalitions (i.e. 2^N , where N is the number of players). In practice, however, the nucleolus is usually reached after solving a small number (often one or two) of these linear programs.

For the bandwidth sharing game, the nucleolus turns out to be

$$B_1 = B_2 = B_3 = \frac{3.5}{3} \text{ Mbps}$$

So in this case the nucleolus is the same as the Nash bargaining solution. It is surprising that even though D1 holds a “veto power”, the nucleolus gives an equal share of the bandwidth to all the providers. This is only true for this example because no two providers can obtain a significant amount of bandwidth by excluding the third provider. For example, if the link between D and E were owned by D2 instead of by D3, the Nash bargaining solution would not change, but the nucleolus would be

$$\begin{aligned} B_1 &= 1.46 \text{ Mbps} \\ B_2 &= 1.29 \text{ Mbps} \\ B_3 &= 0.75 \text{ Mbps} \end{aligned}$$

For the intelligence sharing example, the nucleolus turns out to be

$$s_1 = 2, \quad s_2 = 4, \quad s_3 = 0$$

One reason we get such pessimistic results in the intelligence sharing example is that our game model is not super-additive. That is, the union of two disjoint coalitions does not necessarily produce a coalition with a value greater than the sum of the values of the original coalitions.

The emptiness of the core in the intelligence sharing example implies that the power structure of the game does not clearly point to a single result. In this case the structure of the negotiation and the attitude of the negotiators (i.e. factors such as risk-aversion) have a strong influence on the outcome. Any detail in the negotiation mechanism that tends to steer the players towards a certain outcome could play a major role in determining the result.

The Shapley value [8] assigns payoffs to players depending on their average marginal contribution to the grand coalition. Therefore, it favors players who play a larger role in the success of the grand coalition over smaller players, and is indicative of each player’s power in the coalition. The Shapley value is the unique solution satisfying the axioms of

- *Anonymity*: The solution is independent of the labeling of the players.
- *Carrier (a.k.a. Dummy player property)*: A player who contributes nothing to any coalition does not get any payoff.
- *Linearity*: For any two games defined over the same player set, the solution of the sum game is the sum of the solutions of the individual games.

Computationally, the Shapley value of a player i is given by

$$\phi_i = \sum_{S \subset N} \frac{(s-1)!(n-s)!}{n!} (v(S) - v(S \setminus \{i\}))$$

where N is the grand coalition, n is the number of players in N , s is the number of players in S , and $v(S)$ is the value of coalition S .

The Shapley values for the bandwidth sharing example reflect the imbalance of power among the providers. The Shapley values are

$$B_1 = \frac{17}{12} \text{ Mbps}, \quad B_2 = \frac{14}{12} \text{ Mbps}, \quad B_3 = \frac{11}{12} \text{ Mbps}$$

Thus D1 receives the most bandwidth, because of its vital role, and D3 receives the least due to its small contribution. Also note that the Shapley value lies within the core.

For the intelligence sharing example, the Shapley values are

$$V_1 = 9.91, \quad V_2 = 0.55, \quad V_3 = 10.84$$

This bears out our intuition that in negotiations, A2 is in a weak position because it has less sources than the others, and of less reliability. However, due to the emptiness of the core, the Shapley values are somewhat less informative than in the bandwidth sharing example - they give us information about the distribution of power, but the actual outcome of a negotiation still depends on the details of negotiation.

6. Conclusion

The outcome of a policy negotiation between autonomous domains is often predictable, given some knowledge of the power structure and the criteria used by negotiators to evaluate potential outcomes. The techniques used in such predictions can also be used in automated negotiation agents, reducing the time required to set up dynamic

coalitions for online collaboration. Even if total automation is not desirable, these techniques can be useful as decision-making aids for human administrators in finding a shared access control policy. Game theory provides valuable mathematical tools for such applications.

However, it is clear from our intelligence sharing example that building game theoretic models of negotiations is not always straightforward. The resulting models may turn out to have undesirable properties, such as empty cores. It is not yet clear what models are appropriate for more realistic applications.

All the models in this paper assumed perfect knowledge - all players knew the resources and reward functions of all other players. Modeling of negotiations where each negotiator has only partial knowledge of other negotiators' resources and reward functions is an interesting area for future research, as is the design of effective negotiating mechanisms and strategies in this case.

Acknowledgments

This research effort was sponsored by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-00-2-0510. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, expressed or implied, of the Defense Advanced Research Projects Agency (DARPA), the Air Force Research Laboratory, or the U.S. Government.

References

- [1] D. Shands, R. Yee, J. Jacobs, and E. J. Sebes. Secure virtual enclaves: Supporting coalition use of distributed application technologies. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, San Diego, California, 3-4 February 2000.
- [2] A. Herzberg, Y. Mass, J. Michaeli, D. Naor, and Y. Ravid. Access control meets public key infrastructure, or: Assigning roles to strangers. In *Proceedings of IEEE Symposium on Security and Privacy*, Oakland, California, May 2000.
- [3] N. Damianou, N. Dulay, E. Lupu, and M. Sloman. The ponder policy specification language. In *Policies for Distributed Systems and Networks: Proceedings of the POLICY 2001 International Workshop*, volume 1995 of *Lecture Notes in Computer Science*, pages 18–38. Springer, 2001.

- [4] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [5] F. Forgo, J. Szep, and F. Szidarovsky. *Introduction to the Theory of Games: Concepts, Methods, Applications*, volume 32 of *Nonconvex Optimization and Its Applications*. Kluwer Academic Publishers, 1999.
- [6] J.F. Nash Jr. Two-person cooperative games. *Econometrica*, 21:128–140, 1953.
- [7] D. Schmeidler. The nucleolus of a characteristic function game. *SIAM Journal on Applied Mathematics*, 17:1163–1170, 1968.
- [8] L. S. Shapley. A value for n-person games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, pages 307–317. Princeton University Press, 1953.